

From wang!elf.wang.com!ucsd.edu!packet-radio-relay Thu Feb 7 17:09:26 1991 remote
from tosspot
Received: by tosspot (1.63/waf)
via UUCP; Sat, 09 Feb 91 11:27:15 EST
for lee
Received: from somewhere by elf.wang.com id aa16058; Thu, 7 Feb 91 17:09:25 GMT
Received: from ucsd.edu by uunet.uu.net (5.61/1.14) with SMTP
id AA14800; Thu, 7 Feb 91 11:23:10 -0500
Received: by ucsd.edu; id AA20926
sendmail 5.64/UCSD-2.1-sun
Thu, 7 Feb 91 04:30:20 -0800 for hpbbrd!db0sao!dg4scv
Received: by ucsd.edu; id AA20911
sendmail 5.64/UCSD-2.1-sun
Thu, 7 Feb 91 04:30:09 -0800 for /usr/lib/sendmail -oc -odb -oQ/var/spool/
lqueue -oi -fpacket-radio-relay packet-radio-list
Message-Id: <9102071230.AA20911@ucsd.edu>
Date: Thu, 7 Feb 91 04:30:03 PST
From: Packet-Radio Mailing List and Newsgroup </dev/null@ucsd.edu>
Reply-To: Packet-Radio@ucsd.edu
Subject: Packet-Radio Digest V91 #36
To: packet-radio@ucsd.edu

Packet-Radio Digest Thu, 7 Feb 91 Volume 91 : Issue 36

Today's Topics:

'To:' field anarchy! (2 msgs)
a few random thoughts about channel access (3 msgs)
Has Part 97 changed THAT much? (was Re: PACKET->Internet Gateway)
Internet->packet Gateway (3 msgs)
Painfully long FTP transfers (2 msgs)

Send Replies or notes for publication to: <Packet-Radio@UCSD.Edu>
Send subscription requests to: <Packet-Radio-REQUEST@UCSD.Edu>
Problems you can't solve otherwise to brian@ucsd.edu.

Archives of past issues of the Packet-Radio Digest are available
(by FTP only) from UCSD.Edu in directory "mailarchives/packet-radio".

We trust that readers are intelligent enough to realize that all text
herein consists of personal comments and does not represent the official
policies or positions of any party. Your mileage may vary. So there.

Date: 6 Feb 91 18:06:42 GMT
From: shelby!paulf%shasta.Stanford.EDU@uunet.uu.net (paulf)
Subject: 'To:' field anarchy!
To: packet-radio@ucsd.edu

In article <1991Feb6.033743.23863@maverick.ksu.ksu.edu> steve@matt.ksu.ksu.edu (Steve Schallehn) writes:

>I feel this sort of 'segmentization' is going to be extremely important
>for message distribution continuing in packet radio.

I couldn't agree more. The thing that differentiates netnews from mailing lists is the existence of outstanding filter tools to get at what you want, and to dump all the chad...

-=Paul Flaherty, N9FZX | Without KILL files,
->paulf@shasta.Stanford.EDU | life itself would be impossible.

Date: 6 Feb 91 17:43:20 GMT
From: netnews.upenn.edu!platypus!bill@rutgers.edu (Bill Gunshannon)
Subject: 'To:' field anarchy!
To: packet-radio@ucsd.edu

In article <1991Feb6.033743.23863@maverick.ksu.ksu.edu>, steve@matt.ksu.ksu.edu (Steve Schallehn) writes:

> In the 9th Amateur Radio Networking Conference (Sorry, I don't remember the
> official name), there was a paper presented on using netnews for packet
> radio.

I suggested that right here in this group over 6 years ago. I even tested the idea of using UUCP 'g' protocol with TNC's. It all worked really well.

Of course, I was told that the BBS concept worked perfectly well and there was no need for anything like NEWS. I think it would have been a lot easier to change things before we had as many BBSes as we now do.

bill KB3YV

--

Bill Gunshannon		If this statement wasn't here,
bill@platypus.uofs.edu		This space would be left intentionally blank

Date: Wed, 6 Feb 91 08:00:56 -0800
From: brian (Brian Kantor)
Subject: a few random thoughts about channel access
To: packet-radio, tcp-group

Whilst having packet-radio nightmares again last night, a couple of

thought about channel access methods came to mind. Here I go, challenging assumptions again.

1. CSMA/CA is what we do to avoid collisions - we watch the channel and wait until it's clear. Most sophisticated people do it with p-Persistence. However, I think that a small variation in pP methods would help throughput.

Simply, most of the pP implementations I've played with ALWAYS use the roll-a-die-in-this-slot method - so when they go to transmit a packet, they could conceivably wait one or more slottimes before they transmit even though the channel is clear. I think that if there is no channel activity present the first time you have a packet ready to transmit, you should simply go for it. If the channel IS active, you wait until it's clear, then you start doing the slot-delays.

This variation has the win that you'll never unnecessarily slot-delay on a clear channel, but you will still gain the pP advantage of avoiding the "lets all jump on the channel now that he's done" syndrome which non-pP channel access methods exhibit. Implicit in this is that the generation of packets ready-to-transmit is somewhat random; with maxframe set to one, even a high-volume data source like a BBS or a sending FTP will exhibit this behaviour, since the next packet is, by definition, not ready to transmit until the previous one has been ack'd.

Problem with this one is implementing it; it probably means firmware changes in everyone's TNC. Arrgh. Only the on-the-bus people can do this in software. Still, if you've got a DRSI card, you might give it a shot and see if it helps.

2. Backoff. Exponentially backing off when you don't get an ACK for a packet has one tacit assumption that may be fatally flawed: that the packet or its ack were lost due to congestion that can be cured by reducing traffic density. Yet on a packet radio network where packets are lost randomly due to non-congestion-related causes (like static crashes and passing Volkswagens), this assumption, if applied purely, leads to backing off a lossy channel when in fact the right thing to do is to be more aggressive! (I recall one FTP session that had backed off to an 8-hour retry time because I'd not limited backoff times and it was a really lossy (50% or more dropped) channel.)

Some technique for noticing the degree of congestion and adjusting the retry time is needed - perhaps something can be done along the lines of noting other traffic on the channel and adjusting the exponent in the backoff formula accordingly. Algorithms which give greater weight to the round trip time of successful (i.e., ack'd) packets are also a good idea for combating the pathological-backoff problem that simpler

methods might generate. Gotta think more on this one.

3. p-Persistence slottimes. I'm wondering if we aren't really using far too short a slottime. On a network which has hidden stations (i.e., 90%+ of all ham packet radio networks), waiting a few milliseconds because your coin didn't come heads-up in the current slot isn't going to help - you're still going to stomp on the hidden station's packet that he began transmitting those few milliseconds ago. It seems to me that you'd want the slottime to be more on the order of the transmission time of the average packet, if indeed not of the transmission time of the MAXIMUM packet. That way, when you toss the coin and lose in this slottime, the other guy gets a clear channel for the whole packet, not just the beginning of it. Implicit in the use of short slottimes is the idea that you'll hear him and back off, which simply isn't the case a really high proportion of the time.

Using slottimes on the order of one to five seconds (for a 1200 bps channel) demands that you use technique #1 above; you'd really NOT want to randomly wait some multiple of seconds on a clear channel. It might be smart to do this dynamically - if you're seeing packets but not their acks, or acks but not the packets they're for, you've got hidden stations and you should be using long slottimes. Otherwise you're just fine with slottimes on the order of the DCD response time of your demodulator. Again, this requires quite a bit of smarts, so the on-the-bus people have an advantage, but the this one CAN be done with KISS implementations, since the computer is getting all the packets and can make the determination as to whether there are hidden stations present or not, and adjust the TNC's slottime parameter accordingly.

Comments?

- Brian

Date: 6 Feb 91 23:48:37 GMT
From: epic!karn@bellcore.com (Phil R. Karn)
Subject: a few random thoughts about channel access
To: packet-radio@ucsd.edu

Brian's first two points (decreasing p only when the channel is busy and limiting backoffs) are well taken. It seems to me that both can be set automatically by estimating the number of active stations on the channel. For example, if you see eight active stations on the channel, then p should be 1/8 and the retransmission backoff should be limited to 8. Note that it's the number of active stations and not the actual amount of traffic that matters.

There are two problems to be solved here. One is estimating the number

of hidden stations on the channel and the other is picking an interval during which a station is considered to be "active". One possible way to find hidden terminals is to watch destination as well as source addresses on the channel.

As far as slot times go, I think it's best to keep them short. There's really no way to detect hidden terminals by just listening to channel activity - you have to interpret it. One way is to watch the control fields in the packets themselves - if you see someone send an I frame, you know that its recipient will be sending an ACK soon, even if you can't hear it. This is the basis of the "prioritized ACK" scheme invented by N7CL. I have devised a more general scheme of my own called CSMA/CA (collision avoidance) that is based on the Apple Localtalk channel access protocol; it was described in last year's ARRL conference proceedings.

Phil

Date: 7 Feb 91 00:52:45 GMT
From: brian@ucsd.edu (Brian Kantor)
Subject: a few random thoughts about channel access
To: packet-radio@ucsd.edu

In article <1991Feb6.184837@epic.bellcore.com> karn@thumper.bellcore.com writes:
>... For example, if you see eight active stations on the
>channel, then p should be 1/8 and the retransmission backoff should be
>limited to 8. Note that it's the number of active stations and not
>the actual amount of traffic that matters.

A quibble: I contend that it's the number of stations ready to transmit, not the number of active stations. Assuming point-to-point communications, which is common, and no simultaneous data exchange in both directions, which is rare, the actual number in your scenario would be 4, leading to a pP of 1/4.

- Brian

Date: 6 Feb 91 00:28:04 GMT
From: sdd.hp.com!elroy.jpl.nasa.gov!turnkey!orchard.la.locus.com!
fafnir.la.locus.com!dana@ucsd.edu (Dana H. Myers)
Subject: Has Part 97 changed THAT much? (was Re: PACKET->Internet Gateway)
To: packet-radio@ucsd.edu

In article <11771@helios.TAMU.EDU> willis@photon.tamu.EDU (Willis Marti) writes:
>-----
>In article <446@ultrix.uhasun.hartford.edu>, jlbloom@uhasun.hartford.edu (Jon

Bloom) writes:

```
>
>|> repeater is explicitly allowed [see 97.205(d)]. In the case of the
>
>|> 97.109(e) allows packet stations operating above 50 MHz to pass third-
>|> party traffic under automatic control, but "The retransmitted messages
>|> must originate at a station that is being locally or remotely controlled."
>|> Even worse, messages originated by non-hams (where the notion of a control
>|> operator can't possibly be stretched to cover the originator) surely come
>|> under the requirements of 97.115(b) which states in part:
>|>
>|> (b) The third party may participate in stating the message where:
>|>     (1) The control operator is present at the control point and is
>|>         continuously *monitoring* and supervising the third party's
```

[remainder deleted]

My copy of Part 97 is in the ARRL "The FCC Rule Book". None of these paragraphs (a) exist or (b) say the same thing. Has Part 97 really changed that much since November 1, 1987?

--

```
* Dana H. Myers KK6JQ      | Views expressed here are *
* (213) 337-5136          | mine and do not necessarily *
* dana@locus.com          | reflect those of my employer *
```

Date: 6 Feb 91 15:36:07 GMT
From: magnus.ircc.ohio-state.edu!zaphod.mps.ohio-state.edu!wuarchive!
cs.utexas.edu!helios!photon!willis@tut.cis.ohio-state.edu (Willis Marti)
Subject: Internet->packet Gateway
To: packet-radio@ucsd.edu

Summary (for those quick with the 'n' key): More comments on Internet<->AMPR connectivity, with quotes from a couple of postings.

(Jon Bloom) writes:

I think it is. It has much the same characteristics as a phone patch, in fact. But it's not quite what I understood that people wanted. To the extent that hams are willing to accept those limitations, it seems like a good approach to me.

(reply)

The 'limitations' mean that someone on the Internet can not *start* a connection/session to AMPR. For hams, there is little limitation. When I finish my 'munged' router, I'll have a way for me to initiate from the

Internet.

Also to clarify, I *don't* see AMPR being used to connect other Internet sites to each other...

(Bruce Walker) writes:

Careful. While it is quite possible to configure a router so that no one can successfully initiate a connection to some or all TCP ports (services), it isn't generally possible to configure a router to not forward packets which look like part of an established connection but might not be. Such bogons would be discarded at their final destination, but if they had already crossed the airwaves, the damage would have been done.

(reply)

Correct on the router capabilities. Incorrect, I believe, on the second part. See other comments below.

(-Sam, WB6RJH) writes:

packet are a bit harder, as I guess that you'd have to make the superuser of a particular machine (as designated by the IP address of a packet) be considered the control operator; you'd want to have control on that machine of just who was able to send IP packets to the Internet->packet gateway, and the gateway would have to restrict the routing of IP packets to radio links to those from an approved list of ham-operated Internet hosts. It looks doable, but probably would be messy to implement.

(reply)

Interesting idea about superuser==control operator. But you can't restrict packets to those hosts with ham owners -- what if the ham initiates the connection? Remember, gateways are (must be) two-way. It doesn't make sense to talk about "Internet->packet" instead of "Internet<->packet".

BTW, if you want to look at non-messy ways to implement some kind of access control, look at cisco, inc.'s router manual.

In article <1991Feb6.091808.25403@news.arc.nasa.gov>, sjogren@tgv.com (Sam Sjogren) writes:

|> In article <27536@ucsd.Edu>, brian@ucsd.Edu (Brian Kantor) writes...

|> ><description of forging mail, encryption, etc included by reference>

|> >

|> >I would hope that it's only necessary to make a good-faith effort to

|> >ensure that the sender is a ham. There is no way to be absolutely sure;

|> >it's only a question of how much effort you have to put forth to keep

|> >the pharisees happy.

|>

|> I'd love to be able to operate on this basis, in general. I'm a

|> big fan of honour systems. However, if the asshole bureaucrats

|> are going to be, well, assholes, it's good to know that you can

|> come up with the technology to allow connectivity to continue

|> despite the legal requirements. We have the technology, let's

|> hope that we're not forced to use it.

|>

(reply)

To quote Brian, "There is no way to be absolutely sure;". There are lots of repeaters out there that can't guarantee non-hams are denied access. And the ones that do restrict in some way are a lot less secure than any scheme proposed so far.

And for both of y'all, remember there are other applications besides email that people want & must be considered in gateway/access design.

Cheers,
Willis Marti

Date: 6 Feb 91 18:28:49 GMT

From: sdd.hp.com!zaphod.mps.ohio-state.edu!rpi!uupsi!cci632!cb@ucsd.edu (Just another hired gun (n2hkd))

Subject: Internet->packet Gateway

To: packet-radio@ucsd.edu

In article <1991Feb6.091808.25403@news.arc.nasa.gov> sjogren@tgv.com writes:
>big fan of honour systems. However, if the asshole bureaucrats

^^^^^^

Sorry, if this looks like I'm picking on someone, but THIS is a prime example of why rec-ham.* can't be routed to the packet system. I have devised ways of automatically handling the list of BAD words and such, but then there's always the doubting Thomas that says it won't happen here.

As in the case of this area, doing something new and experimenting and prototyping, etc will not happen. All of those ideas and the people who have them, don't want to take the risk of being wrong, and therefore rather give lipservice than to attempt to fix the problem.

Those of us who post here, might want to consider keeping some of these newsgroups as to the specification of the FCC, after all I might be one of those automatic stations that is passing the traffic, through the Radio system. It would be quite embarrassing to get a citation from Big brother and not even be able to figure out how you deserved it :-).

As far as passing traffic I would consider having a call sign look up function to match the addressor [and the addressee] for verification and thus putting the burden on the originator. The call sign info is available and should be deemed accurate,

afterall didn't the FCC have something to do with it?
other mail would be considered third part mail and be handled
separately...

yet another thought on this subject...

--

email: cb@cci632.cci.com or cb@cci632 or !rochester!kodak!n2hkd!curtis
Curtis Braun, N2HKD, Computronics, PO Box 1002 Fairport NY, 14450

Date: 6 Feb 91 23:32:56 GMT
From: usc!wuarchive!zaphod.mps.ohio-state.edu!sol.ctr.columbia.edu!
cunixf.cc.columbia.edu!cunixb.cc.columbia.edu!mig@ucsd.edu (Meir)
Subject: Internet->packet Gateway
To: packet-radio@ucsd.edu

In article <1991Feb6.182051.2051@lescsse.uucp> gamorris@lescsse.uucp (Gary A.
Morris) writes:

>In <1991Feb6.002926.23780@news.arc.nasa.gov> sjogren@drago.tgv.com (Sam Sjogren)
writes:

>

>>In article <27441@ucsd.Edu>, brian@ucsd.Edu (Brian Kantor) writes...

>>>They way I plan ... to implement an

>>>internet->packet mail gateway is actually rather simple: ...

>>>... Traffic from the internet to the ham side

>>>is filtered; it must be from a mailbox (i.e., contents of the FROM

>>>line) that is on my list of known hams, which is built by observation

>>>and registration.

>

>>It's terribly trivial to create fake mail. I can send mail to you

>>with just about anything in the From: line, using SMTP over TCP.

>>Perhaps this would be a case where we'd need authenticated mail,

>

>Sounds like overkill to me. Couldn't we just say that any unlicensed

>person who sends fake email is illegally operating a amateur radio

>transmitter without a license?

>--GaryM

>--

>Gary Morris	Internet: lescsse!gamorris@menudo.uh.edu
>Lockheed, Houston, Texas	UUCP: lobster!lescsse!gamorris
>Space Station Freedom	Internet: gmorris@nasamail.nasa.gov
>N5QWC/W5RRR	Phone: +1 713 283 5195

Yes; But the FCC will accept this only if you have put a lock on your
system. Some sort of authentication/verification is needed as well as
reasonable checks for illegal traffic. Otherwise, get ready to read ALL of

the traffic first!

(yes; how many of us lock our cars but not our transmitters?
Maybe this is OK if the room gets locked :-)

```
* * * * * * * * ===== Meir Green
* * * * * * * * ===== mig@cunib.cc.columbia.edu
* * * * * * * * ===== N2JPG
```

Date: 6 Feb 91 03:20:05 GMT
From: sdd.hp.com!samsung!munnari.oz.au!uniwa!vax7!nmurrayr@ucsd.edu
Subject: Painfully long FTP transfers
To: packet-radio@ucsd.edu

I am running NOS version 900828 on an XT clone, and I find that FTP sessions, long Telnet sessions and so on can be so slow that I'm likely to be collecting the old-age pension before they finish.

I tracked down one problem. I was running TNC2 ROM version 1.1.7 in my TNC, and it seems that the KISS defaults in this version were wrong. For example, SLOTTIME was set to 50: presumably meaning 500mS. The KISS v4 source I have used 5 (50 mS). Changing mine to this value meant that I actually transmitted a packet now and then on a mediumly crowded channel (sometimes it would take up to 30 seconds on an uncrowded channel. Is there a good way of determining how SLOTTIME and PERSISTENCE should be set?

That's not the whole of the problem, however. It seems that the recovery timer can take ridiculous values. If something goes wrong (e.g. the receiver misses a packet or the transmitter misses the ACK), it can take ages before the transmitter polls the receiver. I was doing a transfer last night, on a frequency with nobody else around, and I had one timer value of five minutes. This means that absolutely nothing happened for five minutes, and I'd just got a packet from the receiver not long before.

It seems to me that this is *FAR* too long. Have I set up something wrong? Is there a default setting that I've missed? And how are these values determined anyway (I could dive into the sources and find out, but it's easier to ask someone who knows).

Suggestions would be greatly appreciated. You might even save the TCP/IP situation here in Perth!

....Ron

--

```
Internet: Murray_RJ@cc.curtin.edu.au      | "This brain is
ACSnet:  Murray_RJ@cc.cut.oz.au           | intentionally
Bitnet:  Murray_RJ%cc.curtin.edu.au@cunivm.bitnet | left blank"
UUCP   : uunet!munnari.oz!cc.curtin.edu.au!Murray_RJ |
```

Date: 7 Feb 91 00:01:29 GMT
From: epic!karn@bellcore.com (Phil R. Karn)
Subject: Painfully long FTP transfers
To: packet-radio@ucsd.edu

If you're using AX.25 connected mode, try setting "ax25 blimit" to an estimate of the number of active stations on the channel. Set the kiss slottime to the keyup delay of the modem, and set the p value to $256/n$, where n is the estimate of active stations on the channel.

You might also set the ax25 irtt to a closer estimate in order to speed convergence to the true round trip time.

Phil

End of Packet-Radio Digest
